



THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: TREIBACH-HECK, et al. Group Art Unit: 3627
Serial No.: 10/003,341 Examiner: Christopher BUCHANAN
Filed: 31 October 2001
Title: Multi-Party Reporting System and Method
Attorney ref: Call-Tell MP

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RECEIVED
MAY - 5 2004
GROUP 3600

AFFIDAVIT UNDER 37 CFR 1.132

Sir:

I am Boris Weissman, Ph.D., and I reside at 1971 W. Middlefield Rd., Mountain View, CA 94043. In this affidavit I explain why, in my expert opinion, the invention defined in the claims of the application identified above should be allowed over the prior art cited by the Examiner. Before I explain the reasoning behind my opinion, I outline my education, work experience, and other professional qualifications.

I) QUALIFICATIONS

- 1) Education
 - a) Ph.D., Computer Science, University of California at Berkeley, Berkeley, CA, 1998, (Dissertation title: "Active Threads: Towards Efficient Fine-Grained Parallelism in Object-Oriented Systems")
 - b) B.S., Computer Science, University of San Francisco, San Francisco, CA 1993 (Summa Cum Laude, GPA 4.0)

2) Work Experience

a) 2001 – present, Senior Member of Technical Staff for Processor Virtualization, VMware, Inc. Palo Alto, CA.

My work focuses on the core virtualization of the x86 microprocessor, including the IA-32 interrupt subsystem, and I am also involved in the design and implementation of a binary translator, as well as of debugging and profiling tools.

b) 1999 – 2001, Technical Lead, Core Services, MRI, Loudcloud, Inc. Sunnyvale, CA

As only the second software engineer to join Loudcloud, I was instrumental in defining original products. Among my many other responsibilities, the team I led was responsible for the implementation of a monitoring appliance for Internet and network service monitoring. This system included a remotely managed transaction monitoring server and a number of service monitors including HTTP(s), ICMP, and DNS monitors. I also co-architected my.loudcloud.com, a customer-facing Managed Services Provider (MSP) portal, which was ranked #1 in the MSP industry by Tier 1 Research in a competition that included Intel, Digex, Exodus, and 25 others. I was specifically responsible for the design and implementation of a distributed storage architecture for over 100,000 server, network, database, and application statistics collected for managed customer infrastructures. The storage infrastructure has been deployed in all Loudcloud data centers worldwide and implements a real-time reporting and analysis framework for the presentation of collected performance data.

c) 1998 – 1999, Java Runtime Engineer, Java Systems Software, SME, Sun Microsystems Sunnyvale, CA

Among other projects, I worked on the design and implementation of the scalable parallel runtime system for HotSpot Java VM on SPARC/Solaris, which is the current execution engine of Sun JDK and has been licensed and ported to several other platforms including Hewlett Packard, Apple and Linux.

d) 1993 – 1998, Doctoral Research, University of California at Berkeley, Computer Science Division and ICSI, Berkeley, CA

For my thesis, I built Active Threads, an extensible scalable runtime for parallel object-oriented languages and systems (targeted at shared-memory

multiprocessors and clusters of SMPs). I led the ICSI Sather compiler implementation effort (1995-1998) and coordinated the implementation work of the international team of postdoctoral fellows and Ph.D. students. I also built a multithreaded extension of CORBA to investigate locality aspects of a client-server interaction for multiprocessing servers.

e) 1990 – 1993, Research Programmer, Physics Research Laboratory, University of San Francisco (USF), San Francisco, CA

I developed post-flight evaluation software for measuring the radiation environment in space on board of NASA Shuttle missions. The software was used to assess the risk of mutations caused by cosmic rays. My responsibilities included designing and building the core system, communication system (for microscopes and other devices), building a graphical user interface, and training the lab personnel.

3) Teaching

1993 – 1996: I assisted in teaching various graduate and undergraduate Computer Science courses at UC Berkeley and USF, including courses in parallel software, software engineering, and data structures.

4) Selected Publications

a) "Performance Counters and State Sharing Annotations: a Unified Approach to Thread Locality," Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII), October 1998;

b) "High-Performance Thread Migration on Clusters of SMPs," with B. Gomes, Parallel and Distributed Computing Practices (PDCP Vol. 2, #2), Special Issue on High Performance Computing on Clusters, 2000

c) "Active Threads: Enabling Fine-Grained Parallelism in Object-Oriented Languages," with B. Gomes, et al., 1998 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '98), Las Vegas, NV, July 1998, pp. 115-122;

d) Efficient Fine-Grain Thread Migration with Active Threads," with B. Gomes, J. Quittek and M. Holtkamp, Joint 12th International Parallel Processing

Symposium and 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP 1998), Orlando, FL, March 1998, pp. 410-415;

- e) "Towards High-Performance Multithreaded CORBA Servers," with Hans-Arno Jacobsen, 1998 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '98), Las Vegas, NV, July 1998, pp. 1410-1417;
- f) "Parallelizing Connectionist Networks for Pattern Recognition: A Library Approach," with B. Gomes, et al., 1998 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '98), Las Vegas, NV, July 1998, pp. 1357-1364;
- g) "Efficient Extensible Synchronization in Sather," with J. W. Quittek, International Scientific Computing in Object-Oriented Parallel Environments, First International Conference (ISCOPE 97), December 1997, Lecture Notes in Computer Science 1343, Springer-Verlag, 1997. pp. 65-72;
- h) "Safe Sharing of Objects in a High-Performance Parallel Language," with B. Gomes, et al., submitted to the European Conference on Object-Oriented Programming;
- i) "Type-Safety and Overloading in Sather," with B. Gomes, et al., International Computer Science Institute, Technical Report TR-97-055, December 1997; and
- j) "A Language Manual for Sather 1.1, " with B. Gomes, et al., International Computer Science Institute TR 97-037.

5) Inventions

- a) U.S. Patent 6,505,275, "Method and Apparatus for Scalable Memory Efficient Thread-Local Object Allocation," issued January 3, 2003;
- b) U.S. Patent 6,687,904, "Method and Apparatus for a Locking Policy based on per-Object Locking History," issued February 3, 2004;
- c) "Web-based Management System for Hosting Service," pending U.S. patent application;

- d) "System for Propagation of Authentication Information," pending U.S. patent application;
- e) "Method and System for Report Generation of Performance Data," pending U.S. patent application;
- f) "System for Logging on to Servers Through a Portal Computer," pending U.S. patent application; and
- g) "System for Communicating with Servers Using Message Definitions," pending U.S. patent application.

6) Professional activities

I am a referee for papers submitted for publication by various international conferences, including Programming Language Design and Implementation (PLDI), Symposium on Parallel Algorithms and Architectures, Parallel and Distributed Processing Techniques and Applications, and I am a Member of the Association for Computing Machinery (ACM), the ACM Special Interest Group on Programming Languages (SIGPLAN) and the ACM Special Interest Group on Computer Architecture (SIGARCH), as well as ACM Multimedia.

II) Opinion

Having carefully reviewed the "Multi-Party Reporting System and Method" U.S. Patent Application 10/003,341 ("Application") as well as the "Method and System for Direct Payroll Processing," U.S. Patent No. 6,347,306 (Swart), I have come to the conclusion that the two have little in common other than the partial overlap of the invention field. In particular, I find substantial qualitative differences with respect to (1) the invention applicability, (2) information life cycle management, and (3) the use of the underlying computing technology.

1) Applicability

The Application addresses the general problem of filtered information sharing, information life cycle, workflow, and report generation. While both the Application and Swart may be applied to payroll processing, Swart is entirely dedicated to this field. On

the contrary, the Application uses payroll processing as only one example of a possible multitude of uses. See, for example, the following two paragraphs:

[0011] Moreover, parameters such as time and expenses are not the only types of information that frequently need to be reported. Employees or affiliates may need to file reports concerning the progress of projects, or the status of inventory, for example, or vendors may wish to send bills quickly in order to speed up the payment process.

[0012] Furthermore, similar problems occur even in contexts where the reporting information is not in any way associated with the recipient of the report. This will often occur in the case of processing of customer orders.

Swart is narrowly focused on payroll processing and offers (see, for example, col. 12, lines 38-40):

A computer-implemented method for automatically paying employees net pay immediately upon a triggering event indicative of completion of a work segment.

2) Information Life Cycle

The Application offers a method for information storage in a structured format along with a set of applicable filtering rules to provide multiple views into data to different parties:

[0014] Using a form processing module, the center automatically extracts report data from the report and stores the extracted report data in memory, and, for at least one third party, it stores a set of party-specific rules in the memory. The received report is then associated with the corresponding third party. Via a publicly accessing transmission network, the center allows access by the third party to at least a selected portion of the extracted report data according to that third party's corresponding party-specific rules.

[0032] Extracted data can then be transmitted (according to a known schedule or in response to a request) to the other recipient parties in any conventional manner, for example in the form of a summary, a detailed report, one or more electronic files in a chosen format, etc. or any combination of such methods.

[0066] Each sender, agency and client is preferably assigned a password and has a privilege level. In general, for example, an agency will be allowed full access to all data relating to or sent by

its contractors, as well as full rights to make annotations. Clients, on the other hand, may be restricted to viewing and annotating forms sent by contractors it has hired. Contractors may be given the privilege to make changes to submitted forms, in which case the form processing system preferably flags for the agency that a change has been made or attempted.

Thus, the Application introduces a full privilege-based information broker that makes certain parts of information available to certain parties. The parties may be limited not only with respect to the scope of information available to them, but also with respect to the kinds of actions they can apply to the information (read-only versus read and modify access). Finally, the system orchestrates the workflow by propagating changes to the documents to various interested parties.

On the contrary, *Swart* offers a simple *linear* information flow model in order to expedite payroll processing without any attempt of information filtering for various involved parties; see col. 2, lines 40, to col. 3, line 3:

One aspect of the present invention provides a computer-implemented model for automatically paying employees net pay immediately upon completion of a work segment. The method includes providing an employee data computer system, a banking computer system and a payroll processing computer system, each interconnected via a computer network. According to this method, time and attendance information is acquired in real-time by the employee data computer system from each employee, including shift completion information indicative of the completion of at least one work segment...

Shift completion information, pay amount information and pay deduction information is provided to the payroll processing computer system via the computer network. The method then calculates, the payroll processing computer system, net work segment pay for each work segment completed by the employee, in real-time and immediately upon entry of the shift completion information.. Finally, data indicative of the net work segment pay is provided to the banking computer system via the computer network upon the completion net work segment pay calculation step to provide the employee with immediate access to the calculated net work segment pay, preferably via direct deposit to the employee's bank account.

Swart offers a real-time *fully automated* solution for payroll calculation and pay distribution. The information flow is, however, strictly linear:

employee data computer → payroll computer → bank computer

Manual intervention is intentionally eliminated from the process in order to optimize the system for expeditious fund distribution. In contrast, the Application provides for an interactive and iterative process with various views of the information available to various parties for a prolonged time period. The data can be updated multiple times by different parties and fund distribution does not necessarily need to take place. Flexibility of information distribution and sharing is the key concept for the Application; *Swart* is concerned solely with the real-time fund disbursement.

3) Technology

The Application and *Swart* employ different technologies to achieve their respective goals. *Swart* is concerned with automatically paying employees immediately upon completion of a work segment and relies on a computer network with the following components: employee computer, payroll computer, bank computer; thus,

Col 2, lines 43-46:

This method includes providing an employee data computer system, a banking computer system and a payroll processing computer system, each interconnected via a computer network.

and

Col 2, lines 54-68:

Shift completion information, pay amount information and pay deduction information is provided to the payroll processing computer system via the computer network. The method then calculates, by the payroll processing computer system, net work segment pay for each work segment completed by the employee, in real-time and immediately upon entry of the shift completion information, by determining (i) gross work segment pay based on the shift completion information and pay amount information and (ii) the net work segment pay by applying the pay deduction information to the gross work segment pay. Finally, data indicative of the net work segment pay is provided to the banking computer system via the computer network upon the completion net work segment pay calculation step to provide the employee with immediate access to the calculated net work segment pay, preferably via direct deposit to the employee's bank account.

Moreover, Swart explicitly states (col. 3, lines 26-31) that the employee data computer is connected to the payroll computer via the Internet:

The system further includes a payroll client computer with at least one payroll software object comprising net pay calculation logic for computing net work segment pay and a payroll application server computer connected to the employee data client computer system and the payroll client computer via the Internet using a programming language-neutral communications protocol.

On the contrary, the Application does not require the Internet to communicate data to the server and explicitly states that such a requirement would be too restrictive:

Paragraph [0004]

One way to reduce the need for data entry would be to have the employee enter data himself into the tracking system. The advantages of this are, however, plain. First, the employee may not be comfortable with or competent in using the time-tracking software, especially if the software is based on a web browser. Second, even assuming that every reporting employee has convenient access to a suitable computer terminal, it would be wasteful to load multiple copies of the time-tracking software into each terminal.

While the application does not dismiss a possibility of completely automatic electronic submission of data, it goes into great detail describing alternate fax and OCR based approaches:

Paragraph [0019]

Reports may be purely electronic, such as an on-line browser input screen, or physical, such as paper . In the cases in which the report is a physical form, the central server includes devices and software modules, in particular, a form-processing application, for automatically receiving an electronic representation of an image of a physical form from any of a plurality of the senders, via a transmission channel.

Paragraph [0027]

[In the cases where the form is paper, the transmission device 120 is preferably a conventional facsimile ("fax") machine. There are several advantages to using a fax machine. First, whereas not all workplaces have convenient Internet connections, almost all have at least one fax machine. Second, almost every sender in the

modern work environment will be familiar with the use and operation of a fax machine. Third, the channel 3000 used by a fax machine is usually the public telephone network, which allows not only for a dedicated, point-to-point connection with the center 400m but also is generally more reliable than a non-dedicated network such as the Internet.

The Application further describes a method for extracting and parsing the data contained in preferably faxed documents for subsequent processing:

Paragraph [0041]

The actual data contained in the different data fields can then be extracted, also using known technology, in a data extraction submodule 433. Several methods for data extraction are known as "optical character recognition" (OCR), which typically recognizes alphanumeric symbols and convert them into corresponding ASCII (or equivalent) character sets; "intelligent character recognition" (ICR), which typically also attempts to recognize non-standard fonts and even handwriting; "optical mark recognition" (OMR), which determines whether a data field (such as a check box, rectangle, circle, etc.) is filled in or blank, and various combinations and variations of these basic concepts.

Swart uses a conventional three-tier architecture, which is common among purely Internet based solutions; see col. 8, lines 41-44:

The present invention preferably uses a three-tier client/server solution. In the three-tier solution used by the present invention, CORBA client JavaBeans "talk to" or communicate with CORBA server objects.

4) Conclusion

I find the Application and *Swart* not only to be substantially different, but in many respects taking diametrically opposing views: *Swart* pursues full automation for real-time processing of pay stubs, whereas the Application provides a method for an interactive analysis of information with many parties involved for potentially substantial time periods. *Swart* requires the use of the Internet and is based entirely on the three-tier Web architecture. The Application prefers to work with physical medium such as paper and may employ fax for communications with a subsequent stage of information extraction. Finally, *Swart* is concerned entirely with payroll processing while the Application provides a general method that is applicable to a multitude of areas.

5) Declaration/Oath

I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

04-28-04

Date

Boris Weissman

Boris Weissman